A Neural Information Field Approach to Computational Cognition

Chris Eliasmith
UNIVERSITY OF WATERLOO

11/18/2016
Final Report

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services, Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
13-03-2017	Final	01 Sep 2013 to 31 May 2017

4. TITLE AND SUBTITLE
A Neural Information Field Approach to Computational Cognition

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA8655-13-1-3084

5c. PROGRAM ELEMENT NUMBER
61102F

6. AUTHOR(S)
Chris Eliasmith, Dimitris Pinotsis

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
UNIVERSITY OF WATERLOO
200 UNIVERSITY AVE W
WATERLOO, 27472356 CA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
EOARD
Unit 4515
APO AE 09421-4515

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/AFOSR IOE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-AFOSR-UK-TR-2017-0016

12. DISTRIBUTION/AVAILABILITY STATEMENT
A DISTRIBUTION UNLIMITED: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The two main research objectives for this project are to understand how local cortical circuits embody states underlying cognitive functions, and to build large-scale models that are able to simulate a variety of tasks and levels of detail. Over the last three years, we have: demonstrated the first large-scale, function neural simulation to use highly detailed single neuron models, allowing the simulated testing of drug effects on cognitive performance; demonstrated a scalable neural model of motor planning; developed a new perceptual decision making model; demonstrated adaptive motor control in a large-scale cognitive simulation with spiking neurons (Spaun); demonstrated simple instruction following in Spaun; shown the first human-scale concept representations in spiking networks; demonstrated how to learn those representations; optimized cognitive computations for the Nengo simulation environment; demonstrated transfer learning to replicate performance of children learning to count in a SPA model; proposed a new SPA model of cognitive load using the N-back task; developed a new model of the effects of distraction in working memory; shown a hippocampal model able to perform context sensitive sequence encoding and retrieval; proposed what is currently the best model of neural activity during context-based working memory retrieval in monkeys; developed new software infrastructure for large spiking neural models; developed specialized hardware implementations of the N-back task; and optimized large model simulations for CPUs.

15. SUBJECT TERMS
EOARD, Computational Cognition, Mixed-initiative decision making

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			LAWTON, JAMES
Unclassified	Unclassified	Unclassified	SAR	11	19b. TELEPHONE NUMBER *(Include area code)* 703-696-5999

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

AFOSR Final Report For: A Neural Information Field Approach to Computational Cognition (FA8655-13-1-3019)

PI: Dr. Chris Eliasmith, University of Waterloo
Co-PI: Dr. Dimitris Pinotsis, University College London

Grant: Sep 1, 2013 to Aug 31, 2016
Submitted: Nov 18, 2016

Abstract

The two main research objectives for this project are to understand how local cortical circuits embody states underlying cognitive functions, and to build large-scale models that are able to simulate a variety of tasks and levels of detail. Over the last three years, we have: demonstrated the first large-scale, function neural simulation to use highly detailed single neuron models, allowing the simulated testing of drug effects on cognitive performance; demonstrated a scalable neural model of motor planning; developed a new perceptual decision making model; demonstrated adaptive motor control in a large-scale cognitive simulation with spiking neurons (Spaun); demonstrated simple instruction following in Spaun; shown the first human-scale concept representations in spiking networks; demonstrated how to learn those representations; optimized cognitive computations for the Nengo simulation environment; demonstrated transfer learning to replicate performance of children learning to count in a SPA model; proposed a new SPA model of cognitive load using the N-back task; developed a new model of the effects of distraction in working memory; shown a hippocampal model able to perform context sensitive sequence encoding and retrieval; proposed what is currently the best model of neural activity during context-based working memory retrieval in monkeys; developed new software infrastructure for large spiking neural models; developed specialized hardware implementations of the N-back task; and optimized large model simulations for CPUs.

Introduction

The two main research objectives for this project are to understand how local cortical circuits embody states underlying cognitive functions, and to build large-scale models that are able to simulate a variety of tasks and levels of detail.

The technical approach followed combines three recent methods developed in Dr. Eliasmith's group. The first method is called the Neural Engineering Framework (NEF), and provides a set of three quantified principles for implementing arbitrary dynamical systems in networks of spiking neurons (Eliasmith & Anderson, 2003). This method relies on a characterization of the function that needs to be computed. That characterization is drawn from a proposed neural architecture called the Semantic Pointer Architecture (SPA; Eliasmith, 2013). The SPA suggests specific functional components and methods for communication critical to mammalian brain function. The SPA provides a general method for constructing neuro-cognitive models that can be implemented by the

NEF. Finally, the Nengo software platform provides a convenient tool for building such models. Nengo is a cross-hardware and OS platform environment for large-scale neural simulation. It can be used as a generic simulator but also incorporates the NEF and SPA approaches as needed.

The expected DOD relevant benefits from this research come from four areas. The first is understanding human processing in dynamically challenging environment. Because the models are all described at the neural level, there is an explicit inclusion of temporal constraints in the models. This leads to the second, more specific benefit, which is better understanding rapid decisions under tight timing constraints. In particular our focus on decision making and context effects in cognition in rapid cognitive tasks addresses this benefit. Third, this research provides a better understanding of decision making and working memory constraints under cognitive load. The N-back task, distraction tasks, and noisy perceptual decision making task all speak to this benefit. Taken together, these insights provide foundations for enhanced performance monitoring. Because the models are built at the neural level, they predict a wide variety of physiological and behavioral measures that can be exploited to better understand human cognitive processing in a variety of challenging environments.

This report covers the research done in Dr. Eliasmith's lab, information on Dr. Pinotsis' research achievements may be requested separately. In what follows, we describe the project goals, describe the central achievements of the project and then conclude with a listing of awards and papers that have been attributed to the grant.

Project Goals

The main goal of the project is to develop a theoretical approach for building *large-scale*, *mechanistic* neural models that match known spatiotemporal properties of cortical dynamics, while realizing specific cognitive functions. In order to achieve this core goal, the research has focused on two specific subgoals.

First, we have identified several specific cognitive tasks or abilities that permit the construction and demonstration of large-scale models, with a focus on comparison to experimental (neural and behavioral) data. In particular, we have focused on decision making, cognitive, and memory tasks. For decision making, we have focused on perceptual decision making, motor planning, instruction following, and extensions to the Spaun model. Our cognitive models have addressed large-scale concept representation, learning of such representations, and learning for task transfer. Our memory models have addressed cognitive load, sequence encoding and retrieval, distraction during retrieval and context based retrieval. Together, these models have allowed us to address general issues of large-scale neural modeling, context dependent cognition, working memory constraints, and learning in cognitive tasks.

Our second subgoal has been to develop theory and infrastructure that allows us to combine functional, physiological, and anatomical neural properties in a single model.

These developments have resulted in new methods for optimizing neural representations for cognitive tasks, methods for introducing high-detail neuron models[1] into cognitive models, neuromorphic hardware running cognitive models, and a large-scale simulation environment (Nengo) that runs large-scale models 450x faster than at the beginning of the project. These methods have been variously used to build the models identified in the first subgoal.

Description of Achievements

In addition to supporting Dr. Eliasmith and Dr. Pinotsis, the project supported several Masters and PhD students at various stages of their career, including Eric Crawford (Masters), Oliver Trujillo (Masters), Peter Duggins (Masters), Sugandha Sharma (Masters), Aaron Voelker (PhD), Brent Komer (PhD), Xuan Choo (PhD), Peter Blouw (PhD), and Jan Gosmann (PhD). The projects that these researchers pursued are described in more detail below. They have been broken down into four relevant areas: decision making, cognitive representation and learning, working memory, and infrastructure.

Decision Making

Specific models that address decision making in the project are quite varied, ranging from low-level perceptual decisions and motor decisions, to high-level action planning and instruction following. These models include (the last names of core personnel are in brackets):

1. NEF model of perceptual decision making (Trujillo)
2. Large-scale motor planning (Blouw)
3. Cognitive behaviour with complex cellular models (Gosmann & Choo)
4. Motor control with adaptation in Spaun (Choo)
5. Instruction following in Spaun (Choo)

The work on perceptual decision making models the famous experiments from the Newsome lab that show monkeys randomly moving dots with varying coherence. Monkeys must indicate which direction the majority of dots are moving. Difficult can be changed by changing what percentage of dots move together. The model was able to reproduce details of the dynamics spiking neural activity, while accounting for the temporal properties of the monkey behavior.

The large-scale planning work determined the set of actions needed to be performed to accomplish a goal, such as 'boil the kettle'. The planning was fully within the SPA spiking architecture and shown to be able to scale extremely well, with 25,000 actions,

---

[1] Any model constructed in the Neuron simulator can be included in these models. We have demonstrated models with 10 ion channels, 15 compartments, and described by about 100 equations. These models are of similar complexity to those in the Human Brain Project.

hundreds of goals, and hundreds of locations, planning was very robust, with both high (>94%) success rates and resilience to error.

The Spaun model was extended to 'BioSpaun', integrating detailed single neuron conductance models in frontal cortex. The effects of TTX on multiple decision making tasks was determined, demonstrating an integrated modeling approach from high biological detail to cognitive tasks within Spaun. This work is being prepared for publication at a high impact venue.

In the adaptive motor control work, decisions regarding movements were updated based on the existence of an unknown force field. This was integrated into Spaun and used to demonstrate that Spaun could adapt to such a force field while writing digits. Furthermore digits not written during training were also adapted for, demonstrating the excellent generalization of the learning.

Finally, preliminary results on instruction following in Spaun were completed as part of the project. This is a complex decision making task requiring the system to parse input, determine what actions to take based on the input, and then execute those actions at the appropriate moments. While preliminary, this work sets the foundations for construction of a much more capable and flexible agent using the SPA.

Cognitive Representation and Learning

Research on cognitive representation and learning focused on the general issues of large-scale representation of concept spaces, efficient neural representation of such spaces, and learning both such representations and their deployment. These models include:

1. Human-scale concept representation in spiking networks (Crawford)
2. Learning human-scale cognitive representations (Crawford & Voelker)
3. Optimizing cognitive computations (Gosmann)
4. Learning in a large-scale cognitive model (Aubin)

The human-scale concept representation research used WordNet, which contains 117K concepts, hierarchically related by 5 different relations (e.g. contains, is a, etc.). This conceptual structure was encoded into an SPA network, and then used to perform variety of tasks. Specifically, the ability to retrieve the information, answer questions regarding the relational relationships (up to 13 hierarchical connections away), and use the representations to encode sentences with embedded structure were tested. All results demonstrated 95% or higher accuracy in a 2.5 million neuron spiking network.

Additional mechanisms were developed in the form of a combined supervised and unsupervised learning rule to allow these same kinds of representations to be learned from data. Initial proofs of concept were performed that were able to learn representations with equal dimensionality to those used the WordNet network. Extensions of this work to a full conceptual structure is ongoing.

When representing high-dimensional concepts in a spiking network, the accuracy is important for supporting effective performance across tasks. This work focused on optimizing the sampling method for generating these high-dimensional representations in order to minimize the number of spiking neurons that were needed for a given decoding accuracy.

One task demonstrated by Spaun was counting or addition (e.g. Given '3' what do you get when you count '4'? Answer: 7). Children are known to perform addition starting with a counting strategy, but eventually transition to a memorization strategy. We built a network to show how cortex could learn to take over a slower, basal ganglia driven counting strategy through memorization. This example of automatization of a task was performed within a standard SPA architecture.

Memory

The models of working memory that were constructed during this research focus on the effects of cognitive load, context, distraction, and sequences. These models include:

1. Modeling cognitive load in working memory (Gosmann)
2. Context sensitive sequence encoding and retrieval (Trujillo)
3. Multi-scale cognitive behavior modeling (Duggins, Gosmann & Choo)
4. Context-based working memory retrieval (Komer, Sharma)
5. Distraction in working memory retrieval (Gosmann)

The cognitive load task we focused on is the well-known N-back task, which has been modeled by several groups. We demonstrated a SPA model that matched experimental behavioral data in detail, and predicted a specific difference in performance between match and mismatch trials that remains to be tested.

One of the central working memory structures is the hippocampus. We built a novel model of the hippocampus that encompasses the main anatomical areas of the hippocampal complex, while reproducing a wide variety of hippocampal dynamics (e.g. phase procession, theta waves), and function. We focused on learning novel sequences and retrieval of those sequences, compare neural activity in the model to recent results in neuroscience that identified 'time cells' that were critical for such functions.

We developed and employed new methods for modeling a working memory tasks across levels of detail, and demonstrating different methods for predicting drug effects (specifically phenylephrine and guanfacine) at the neural and behavioural level. These methods included functional, neural and biophysical manipulations that were able to account for the observed behavioural changes.

Recent work in the Newsome lab has extended the standard random dot task (see above under Decision Making) to include a context signal. This signal indicates to the animal that the memory task will focus on motion or color responses. We have extended our perceptual decision making model to account for the effects of context in this flexible

working memory task. We have shown extremely good fits to behavioral data under both contexts, as well as excellent characterization of complex single cell dynamics during the task. This model clearly outperformed past models in explaining the details of the neural dynamics.

We have constructed a model of context encoding in working memory that is able to reproduce a variety of experiments testing the effects of distraction during list memory. These distractions include short and long delays before recall, and continuous distraction (forced rehearsal of irrelevant information) during the task. The spiking neural model accounts for the probability of first recall, recency effects, primacy effects, and the conditional response probabilities as a function of item lag. This model is being extended to incorporate the above hippocampal model, to provide a full understanding of the working memory system.

Infrastructure

All of the above research relied on significant advances in the infrastructure available for building large-scale cognitive models at the level of spiking neurons. These infrastructure improvements include:

1. Software infrastructure for large models (Crawford)
2. Hardware implementation of cognitive computations (Gosmann)
3. Optimizing large model simulations (Gosmann)

We have extended the Nengo simulation environment by building a new backend for a high performance super computer. This backend allows simulation of any Nengo defined model in an MPI environment, supported by most large parallel computers. We demonstrated excellent scaling and simulations over two orders of magnitude faster than on state-of-the-art workstations.

We also independently constructed a backend that allows simulation of arbitrary Nengo models on SpiNNaker neuromorphic hardware. We used the N-back task to test and improve the hardware mapping, which allows all models to run in real-time.

In order to efficiently run Nengo models when a super computer is not available, we developed methods for optimizing the compute graph that provides a 4x improvement in runtime on standard CPUs. This is critical during model development, debugging, and testing, as interactions with super computer resources is often time consuming (e.g. queuing jobs, waiting for the resource to become free, etc.).

Conclusion

Overall, the research has provided significant progress addressing each of the core goals of the initial proposal. Not only has infrastructure been developed to provide improvements in computational efficiency, the majority of effort has gone to exploit those resources to answer a broad variety of questions about working memory, decision

making, and cognition across the neural and behavioral levels. Models of concept representation and whole brain function have been described at internationally unprecedented scales, and with a greater degree of neural detail than has been accomplished before. This has provided the ability to examine physiological interventions (e.g. drug application), characterize transfer learning, and build significant extensions to the world's largest functional brain model, Spaun. We look forward to pursuing the many avenues of research significantly accelerated by this project.

Awards and Recognitions

Descriptions of the awards are slightly modified from those on the respective websites.

Allen Newell Award ICCM – P. Duggins, T.C. Stewart, X.Choo, and C. Eliasmith. Effects of guanfacine and phenylephrine on a spiking neuron model of working memory. In ICCM. 2016.

The Allen Newell Award – named in recognition of one of the founders of the field of cognitive modeling – is given for the best full paper with a student as first author. The winner was selected by an international jury that judge papers that receive an outstanding review during peer review for acceptance to the conference.

Best Higher-level Cognition Model Prize Cognitive Science Conference – S. Aubin, A. R. Voelker, and C. Eliasmith. Improving with practice: a neural model of mathematical development. In 38[th] Cognitive Science Society, 2012–2026. 2016.

Four prizes are awarded for the best full paper submissions that involve computational cognitive modeling. The four separate prizes represent the best modeling work in the respective areas of: perception/action, language, higher-level cognition, and applied cognition. Each prize includes a cash award of $1,000 (USD). The prizes are open to researchers at any level (student, post-doctoral fellow, research scientist, faculty) from any nationality. Any form of computational cognitive modeling relevant to cognitive science is eligible.

NSERC Polanyi Prize – C. Eliasmith, 2015. http://www.nserc-crsng.gc.ca/Prizes-Prix/Polanyi-Polanyi/Profiles-Profils/Eliasmith-Eliasmith_eng.asp

The award was established in tribute to the excellence in research that John C. Polanyi continues to exemplify. Dr. Polanyi won the 1986 Nobel Prize in Chemistry. NSERC is proud to offer a prestigious research prize in his name.

Created in 2006, the NSERC John C. Polanyi Award is given to an individual or team whose research, conducted in Canada, has led to a recent outstanding advance in any NSERC-supported field of the natural sciences or engineering. The research leading to the advance must have been funded at least partially by an NSERC grant. The award is open to all researchers, regardless of their career stage.

As part of the award, the winning individual or team receives a research grant of $250,000.

Papers Attributed to the Grant

J. Gossman and C. Eliasmith. Optimizing Semantic Pointer Representations for Symbol-like Processing in Spiking Neural Networks. *PLoS ONE*. 2016.

P. Duggins, T.C. Stewart, X.Choo, and C. Eliasmith. Effects of guanfacine and phenylephrine on a spiking neuron model of working memory. In ICCM. 2016.

S. Aubin, A. R. Voelker, and C. Eliasmith. Improving with practice: a neural model of mathematical development. In 38th Cognitive Science Society, 2012–2026. 2016.

I. Kajić, J. Gosmann, T.C. Stewart, T. Wennekers, and C. Eliasmith. Towards a cognitively realistic representation of word associations. In 38th Cognitive Science Society, 2183–2188. 2016.

S. Sharma, B. J. Komer, T.C. Stewart, and C.Eliasmith. A neural model of context dependent decision making in the prefrontal cortex. In 38th Cognitive Science Society, 1122–1127. 2016.

S. Sharma, S. Aubin, and C. Eliasmith. Large-scale cognitive model design using the Nengo neural simulator. *Biologically Inspired Cognitive Architectures*. 2016.

P. Blouw, C.Eliasmith, and B. Tripp. A scaleable spiking neural model of action planning. In 38th Cognitive Science Society, 1583–1588. 2016.

E. Crawford, M. Gingerich, and C. Eliasmith. Biologically plausible, human-scale knowledge representation. *Cognitive Science*. DOI: 10.1111/cogs.12261, 2015

O. Trujillo. A spiking neural model of episodic memory encoding and replay in hippocampus. Master's thesis, University of Waterloo, Waterloo, Ontario, 2014.

O. Trujillo and C. Eliasmith. A spiking-neuron model of memory encoding and replay in hippocampus. *Computational Neuroscience Society (CNS)*, p. 166, 2014.

D. A. Pinotsis, Neural Field Coding of Short Term Memory, Cognitive Science Society Conference, 2014.

T. Bekolay, J. Bergstra, E. Hunsberger, T. DeWolf, T. C Stewart, D. Rasmussen, X. Choo, A. R. Voelker, and C. Eliasmith. Nengo: a python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*. 2014.

A. R. Voelker, E. Crawford, and C. Eliasmith. Learning large-scale heteroassociative memories in spiking neurons. In *Unconventional Computation & Natural Computation*. London. 2014.

E. Crawford, M. Gingerich, and C. Eliasmith. Biologically plausible, human-scale knowledge representation. In 35th *Annual Conference of the Cognitive Science Society*, 412-417. 2013.

A. Hurzook, O. Trujillo, and C. Eliasmith. Visual motion processing and perceptual decision making. In 35th *Annual Conference of the Cognitive Science Society*, 2590-2595. Cognitive Science Society, 2013.

C. Eliasmith and O. Trujillo. The use and abuse of large-scale brain models. *Current Opinion in Neurobiology*, 2013.